Cyberinfrastructure in Academia: A Case Study

- Building a New World
- Two Case Studies:
  - A Researcher-Driven Computing Center
  - A "Supercomputer with an Accelerator Running Through It"
- Conclusions

UPRM PDC Workshop
Mayaguez, Puerto Rico
February 10-11, 2004

Paul Sheldon
Vanderbilt University

A Third Discovery Paradigm

Computation complements **Theory** and **Experiment**

Issues for Science and Engineering Researchers in the Digital Age

A report of the National Academy of Sciences (2001)

"…the exploding technology of computers and networks promises profound changes in the fabric or our world.

As seekers of knowledge, researchers will be among those whose lives change the most.

…**Researchers themselves will build this New World largely from the bottom up**, by following their curiosity down the various paths of investigation that the new tools have opened. It is unexplored territory.

…the hoped-for benefits of these systems will depend on their being made available widely and equitably"

How are University Researchers Exploring this New World?

Two Examples from One University

- **This is not you father's University Computer Center:** Cyberinfrastructure as an Investigator Driven Discovery Tool

- **A Supercomputer with an Accelerator Running Through It:** Grid computing and Fault Adaptation in Quasi-Real-Time Systems

Case 1: A New Twist on the "Campus Computer Center"

~4 years ago, a few physicists & biologists asked:

- Can we **agree on hardware**?
- Is there a **sharing mechanism** that can keep us all happy?
- Will our **cultures clash**?
- Will there be any **synergy**?
- Is a **grassroots**, bottom-up effort **sustainable**?

Demonstration Project:

VAnderbilt Multi-Processor Integrated Research Engine

Experiment a Success…

- Our concerns were unfounded
- Increased rate of discovery
- Brought together a diverse community including new investigators
- Enhanced education
- Responsive to investigators
- Helped recruit excellent faculty
- Attracted External Funding

This encouraged us to try the next step…

…and convinced Vanderbilt to give us $8.3M in seed money…
(funding began October 1)

Vanderbilt Scientific Computing Center (VSCC)

An Investigator Driven Discovery Tool

- **Application Driven**: rather than emphasizing the development of computational hardware, tools and methodologies, we emphasize the application of computational resources to important questions in the diverse disciplines of Vanderbilt researchers,
- **Low Barriers**: provide computational services with low barriers to participation, working with researchers to develop and adapt HPC tools to their avenues of inquiry,
- **Expand the Paradigm**: work with members of the Vanderbilt community to find new and innovative ways to use computing in the humanities, arts, and education,
- **Promote Community**: foster an interacting community of researchers and develop a campus culture that promotes and supports the use of HPC tools.

Diverse and Broad Spectrum of Researchers

~100 Investigators, 19 Departments, 4 Schools

VSCC: A Cross-Fertilization Engine Fueling Discovery

- National Supercomputing Centers:
  - **Remote**
  - **High barriers** to participation (especially for novices)
  - **Insufficient resources** for VU researchers
  - **No educational opportunities** for our students
  - Aren't **responsive** to the needs of a **diverse community**
  - Doesn't help **recruit** the **best students** and **faculty**
  - Doesn't produce a **local culture** and **community**
  - Does not **propel VU** to the **front rank** of **US Universities**
- The point of this center is the culture it will establish, the community it will foster, the educational opportunities it will create, and the synergy that will ensue.

Three Kinds of Users

- **Established** High Performance Computing users
- **Novice HPC** Users, experienced w/ Scientific Computing
- **Agnostics** (doubtful or noncommittal)

Multifactor Dimensionality Reduction

An Established User

"New statistical method allows researchers to associate a triple-gene interaction with increased breast cancer risk"

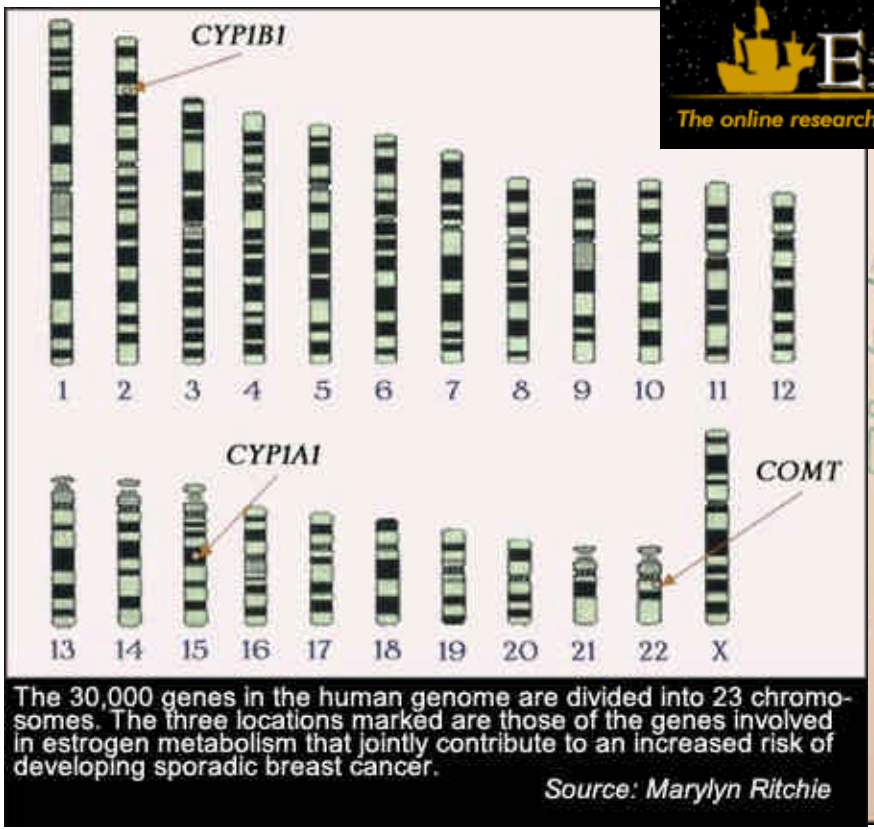

The 30,000 genes in the human genome are divided into 23 chromosomes. The three locations marked are those of the genes involved in estrogen metabolism that jointly contribute to an increased risk of developing sporadic breast cancer.

Source: Marylyn Ritchie

Exploration

The online research journal of Vanderbilt University

The American Journal of Human Genetics

Multifactor Dimensionality Reduction

An Established User

"New statistical method allows researchers to associate a triple-gene interaction with increased breast cancer risk"

The SCC Fosters Cross-Fertilization and Synergy Between Researchers

- Sharing of Data Mining Techniques: first application of Genetic Programming Techniques in Elementary Particle Physics
- Working Together on **National Initiatives** (NSF, DOE) developing **Computational and Data Grid Technology**

Genetic Programming

- A method for optimization.
  - Example: Search for combinations of genes that indicate a clinical outcome [Gene A and Gene B but not Gene C unless Gene D…]
  - Selectively searches a combinatoric space too large to search systematically
- A Population of "programs" is spawned…
  - Programs are made up of functions (mostly operators), variables, and constants
- …The best programs in a population reproduce, yield the next generation
  - Sexual (combine two) and asexual (self-copy) reproduction
  - Mutation
  - Natural selection: "survival of the fittest"
- Successive generations should improve
- Eventual program is transparent (unlike neural net)

First Application of GP in Elementary Particle Physics

- Adaptation of code developed by Human Genetics researchers, worked in concert with them initially
- Evolving program is one that selects candidates for a particular decay process of interest
- Used in searches for extremely rare processes in a very large dataset.
- First indications:
  - GP method can significantly improve background rejection and acceptance of signal (factor of two improvement in significance in at least one case).
  - 30 or so generations typically required
  - Systematic errors understandable (and not significantly larger)
- Publications soon!

Simulations of Devices in a Typical Space Mission

New HPC User Just Coming on Board

Institute for Space and Defense Electronics

U.S. Navy, Draper Lab support (~$2.5M/yr beginning 10/03)

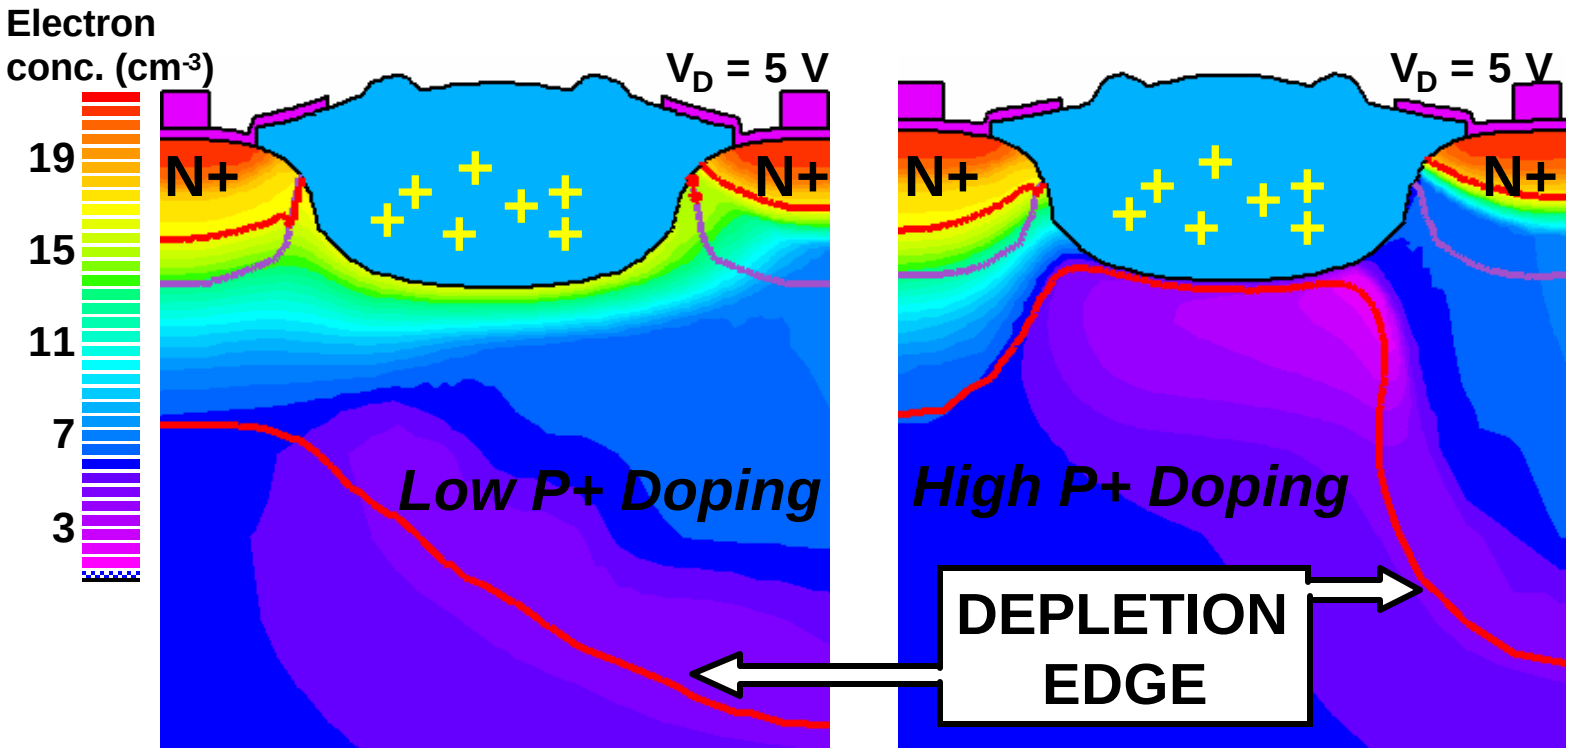

Dose Rate	Applied Bias	Hole Trap Density
0.013 Rad(SiO_2)/s	V_D = 5 V, V_S = V_b = 0 V	$N_T = 10^{17}$ cm^{-3} (Spatially Uniform)

Simulations of Devices in a Typical Space Mission

New HPC User Just Coming on Board

Institute for Space and Defense Electronics

U.S. Navy, Draper Lab support (~$2.5M/yr beginning 10/03)

The VSCC Leverages & Enhances the work of Campus Research Centers

- Last year the Navy invested $250K in VAMPIRE to provide resources for one ISDE research group

Dose

0.013 Rad(SiO_2)/s $V_D = 5$ V, $V_S = V_b = 0$ V $N_T = 10^{17}$ cm^{-3} (Spatially Uniform)

Other Examples of Users…

- Cognition/Neuroscience
- Modeling Supply Chain Management Strategies (Business)
- Supernova Cosmology Project
- Structural Biology (AMBER, …)
- Materials Science
- Many of these users are a new breed…

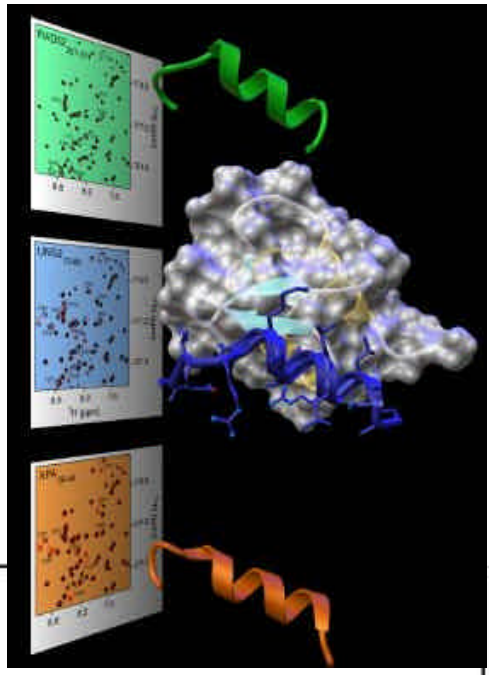

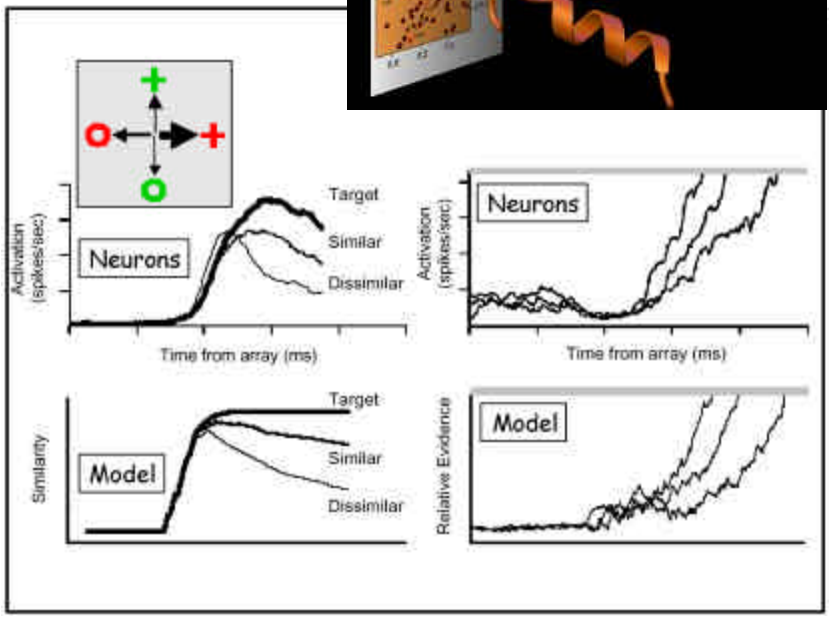

A New Breed of HPC User

- Generating lots of data
  - Some can generate a Terabyte/day
  - No good place currently to store it… (CDs don't cut it)
  - Develop simple analysis models, and then can't go back and re-run when they want to make a change because data is too hard to access, etc.
- These are small, single investigator projects. They don't have the **time**, **inclination**, or **personnel** to devote to figuring out what to do (how to store the data properly, how to build the interface to analyze it multiple times, etc.)
- On the other hand, **money** is not an issue…

User Services Model

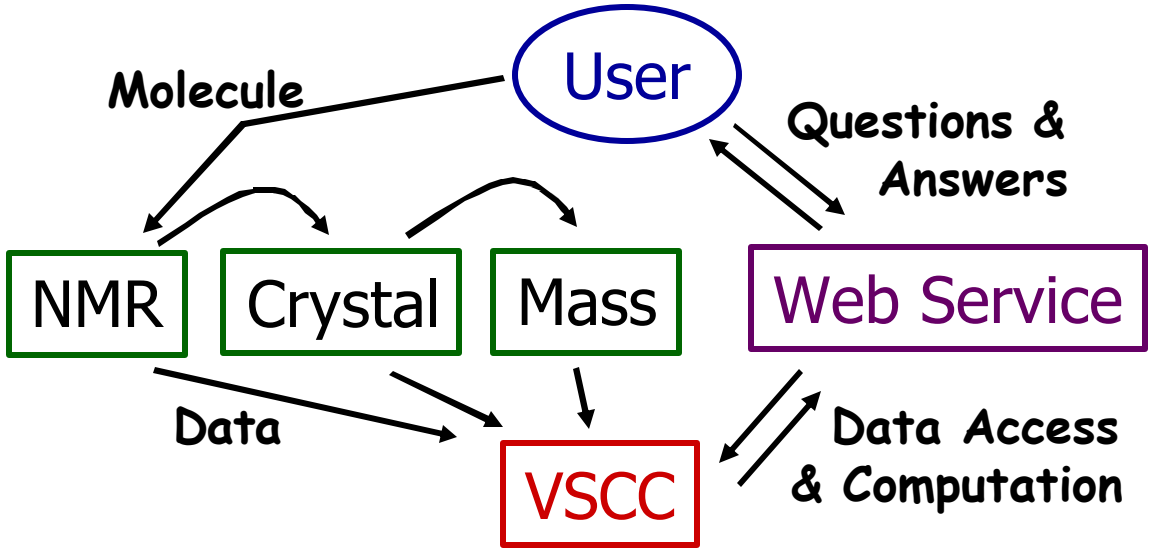

- User has a biological molecule he wants to understand
- Campus "Facilities" will analyze it (NMR, crystallography, mass spectrometer,…)
- Facilities store data at VSCC, give User an "access code"
- Web Service is created to allow user to access and analyze his data, then ask new questions and repeat…

VSCC Components

- Pilot Grants for Hardware and Students
- Educational Program
- Compute Resources
- Storage
  - Tape, low-cost disk, and SAN
- Backup
  - Tape backup and Archive

Pilot Grants & Awards

- 2-year "seed" grants for Vanderbilt faculty ($10K ?$25K)
- ½-time graduate or post-doc support
  - Develop computational expertise within research group
- In addition, for Humanities Faculty:
  - Travel money to present results at conferences
  - Page charges for publications
- Matching funds for external grants.
- Yearly internal competition
- Foster development of:
  - expertise within a research group so can seek external funding
  - new avenues of inquiry in groups w/ minimal/no previous HPC use

Educational Program

- Undergraduate Minor in Scientific Computing
- Graduate Certificate in Scientific Computing
- New courses…

High Performance Computing Course

- Greg Walker (ME) and Alan Tackett (Physics, VSCC)
- Purpose: Apply HPC to actual research projects. Not toy problems.
- Each student is working jointly with a faculty member on a current research project.
- Course Broken into 3 Modules
  - HOW-To's: Makefiles/compiling, cluster design, Parallel Arch, Security
  - Tools: DDT, Dakota, Global Array, PETSc, Matlab, BLAS, FFTW, LAPACK, GSL
  - Programming: MPI, Loosely-coupled vs. Tightly-coupled applications, profiling, parallel debugging, symbolic computing

VSCC Compute Resources

- Eventual cluster size (estimate): 2000 CPUs
- Plan is to purchase ~1/3 of the CPU each year
  - Old hardware removed from cluster when maintenance time/cost exceeds benefit
- 2 types of nodes depending on application:
  - Loosely-coupled: Tasks are inherently single CPU. Just lots of them!
  - Tightly-coupled: Job too large for a single machine. Typically requiring a high-performance networking, such as Myrinet.
- Actual user demand will determine:
  - numbers of CPUs purchased
  - relative fraction of the 2 types (loosely-coupled vs. tightly-coupled)

Diverse Applications

- Serial jobs. But lots of them!
  - High Energy and Nuclear Physics
  - Good for keeping cluster busy
- Small/medium parallel jobs requiring 2-20 CPU's
  - Requires high-performance network
  - Amber (MD, Protein), Human Genetics applications
- Large parallel ASCI jobs using 10-512 CPU's
  - Requires high-performance network
  - Socorro(Condensed Matter Physics)
    - 16 CPU run: 600s with Fast Ethernet vs. 4 sec with Myrinet

Software Libraries

- Because of diverse user group there is a diverse group of software installed
  - Libraries: ATLAS/BLAS, LAPACK, FFTW, PETSc, DAKOTA, Matlab, Netsolve, IBP, MPICH, PVM
  - Compilers: Multiple gcc versions supported, Intel C/C++/F95, Absoft F77/F95
- Users not capable of building these packages. In fact they may not even know they exist!
- Most need to be compiled locally to maximize performance

Resource Sharing: Maui†

- Provides each group on average their appropriate fair share of the cluster
- Supports advanced reservations
  - Serial and parallel jobs
  - Node attributes for special hardware or applications
- Configurable Job priority based on
  - Group, user, account, QoS, number of CPU's, execution time, etc.
- Shortpool Queue for interactive debugging of jobs and short jobs
- Showbf command

†http://www.supercluster.org/

Cluster Building Block

- A Brood is a...
  - Gateway
  - Switch
  - 20 or more compute nodes
- Gateway responsible for
  - Health monitoring
  - Updates and Installs
  - Compute Nodes DHCP service
  - Exporting of /usr/local to nodes
- Brood Flexibility
  - Complete Mini-Cluster
  - Can be segregated from main cluster for users specialized needs.
    - Testing special hardware, kernels, different OS's, apps
  - Easily reintegrated with larger cluster using SystemImager

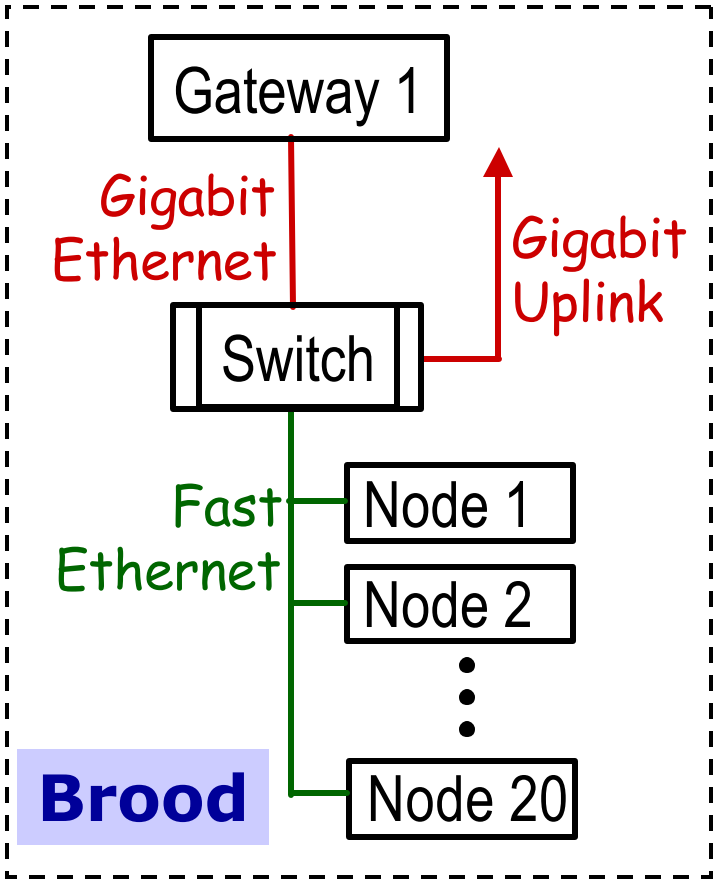

Putting It Together

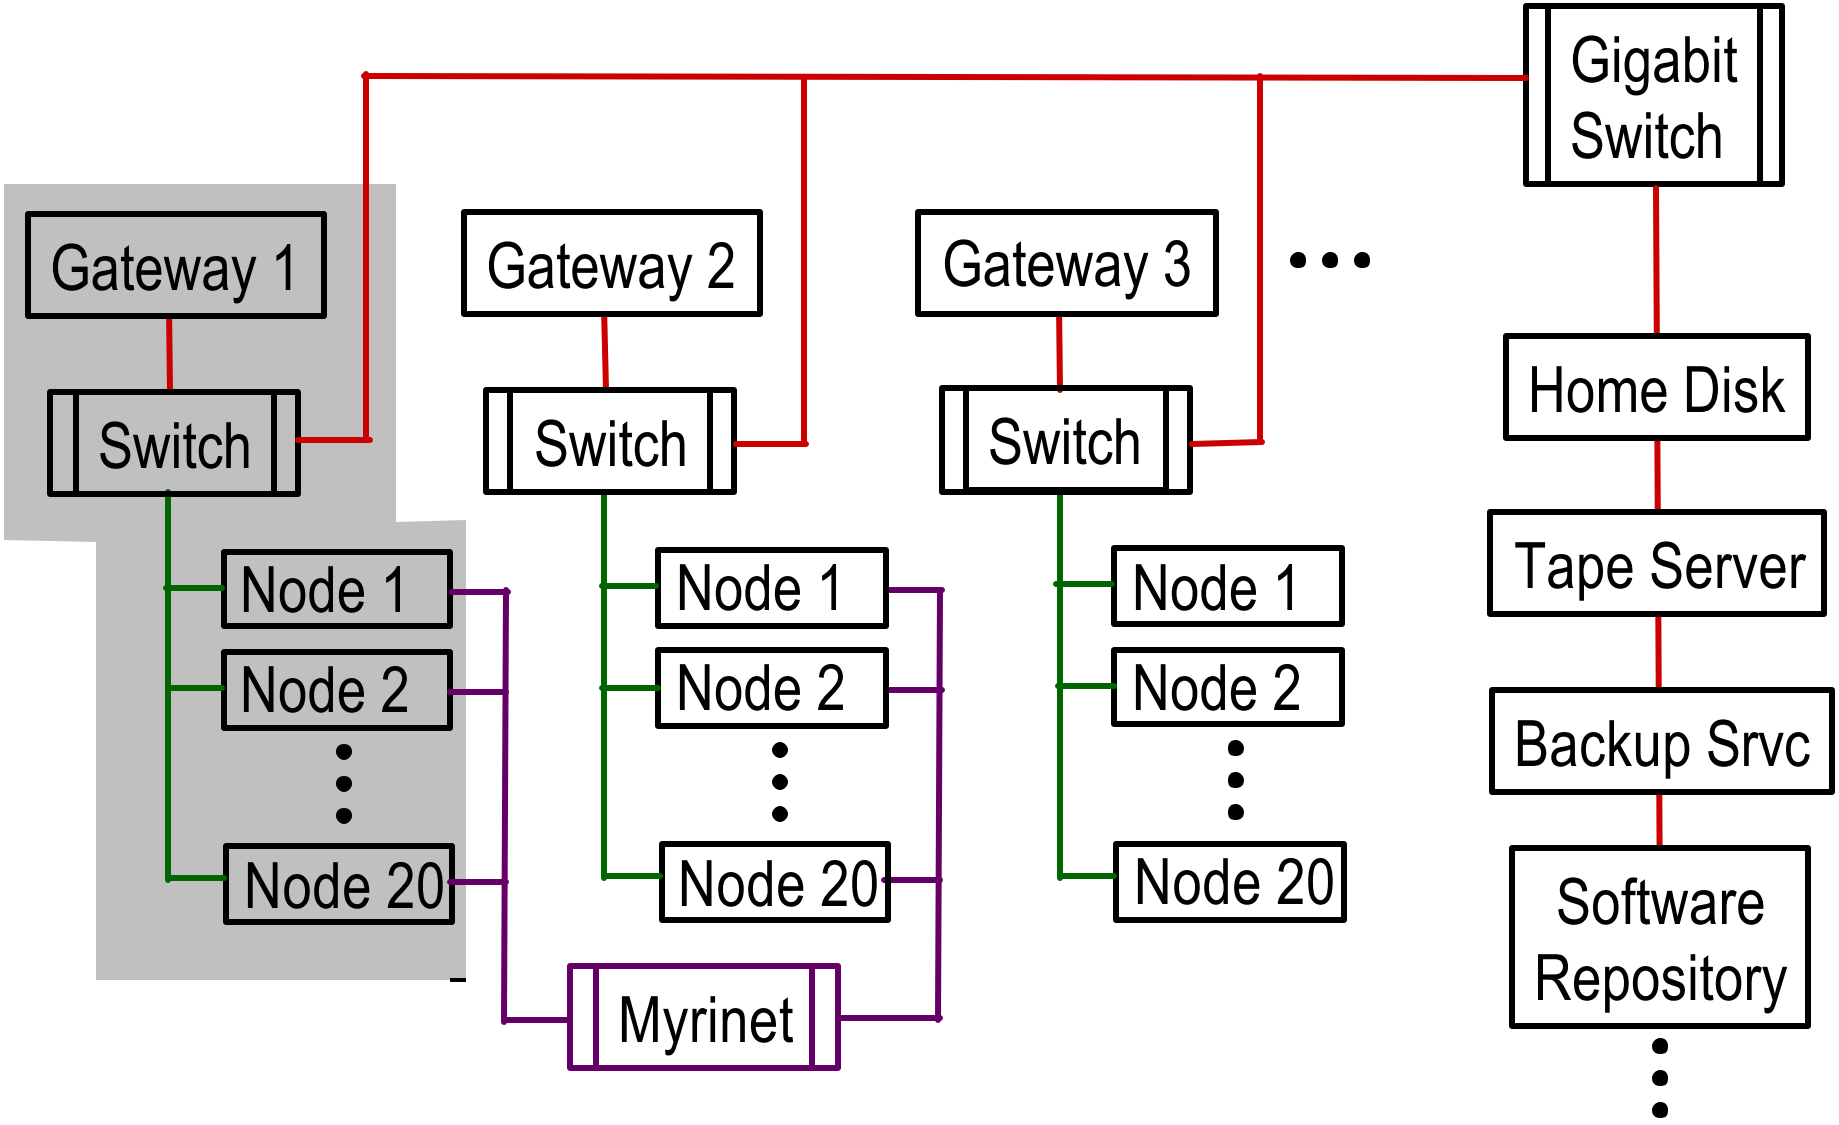

VSCC Economic Model

- Center must be **self sustaining** in 5 years, initial grant is start-up
- Users contribute to the center in any way that they can
  - Some find it easier or only possible to contribute hardware (or personnel).
  - Some prefer to pay users fees, some can't or find it difficult.
  - In kind contributions…
- These must be translated to "Center Dollars" that can be used to "purchase" services
- Example: user buys 30 compute nodes. Price includes support, operations, and maintenance cost. User is "guaranteed" access to those nodes at all times. In addition, they can compete for "excess" CPU cycles that are not currently in use.
- Resources and Center Budget will be determined by the users themselves

Evaluation Metrics

- How will we monitor performance of center and gauge our level of success, both for internal feedback and for reporting to users, university administration?
- Short Term Metrics
  - Number, Diversity of New Faculty & Student Users
  - New inter-departmental and inter-school collaborations
  - Feedback from Users and Investigators
- Long Term Metrics
  - New Faculty SCC Helped Recruit
  - Publications
  - External Funding for Center Researchers
  - Funding for Center
  - External Reviews

Case 2: A Supercomputer w/ an Accelerator Running Through It

- BTeV Experiment has identical computational needs to LHC expts.
- The BTeV Trigger is a "Model Application" for CS researchers investigating high performance, heterogeneous, large scale systems that need to be fault tolerant and fault adaptive…

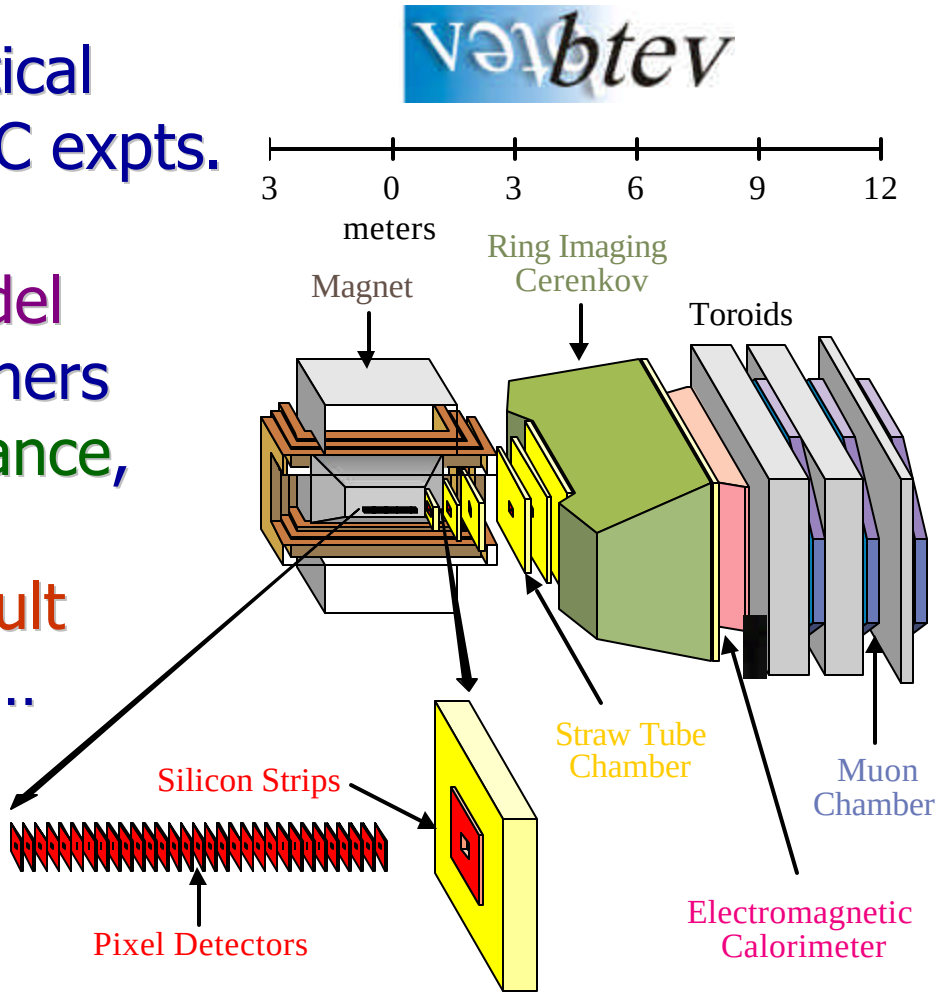

What is BTeV?

- BTeV is an experiment designed to challenge and confront the Standard Model description of CP Violation in Heavy Quark Decay
- Will run at the Fermilab Tevatron, concurrent with LHC.
- Will be the "Flagship Accelerator Experiment" in the US (Mike Witherall)
- Typical HEP "worldwide" collaboration…
  - China
  - Italy
  - Russia
  - US
  - Others…

BTeV is a Petascale Expt.

- Even with sophisticated event selection that uses aggressive technology, BTeV produces a large dataset
- 4 Petabytes of data/year (not that far from ATLAS/CMS)
- Require Petaflops of computing to analyze its data
- Resources and physicists are geographically dispersed (anticipate significant University based resources)
- To maximize the quality and rate of scientific discovery by BTeV physicists, all must have equal ability to access and analyze the experiment's data…
- …sounds like the grid (???)

BTeV Interest in GRIDs

- **Unique Requirements**
  - Dynamic reallocation of grid resources
    - Use Grid Resources (at Universities, say) in online trigger
    - Use Trigger Computing for offline analysis when "idle"
  - Won't use tape: secure widely-distributed disk based data store
- Joined iVDGL, participating in Grid2003 project
  - Vanderbilt node on Grid2003 grid
  - BTeV MC application, full data provenance w/ Chimera
  - VDT Testers
- BTeV Grid Testbed and Working Group forming now

The "Supercomputer – Accelerator" Thing

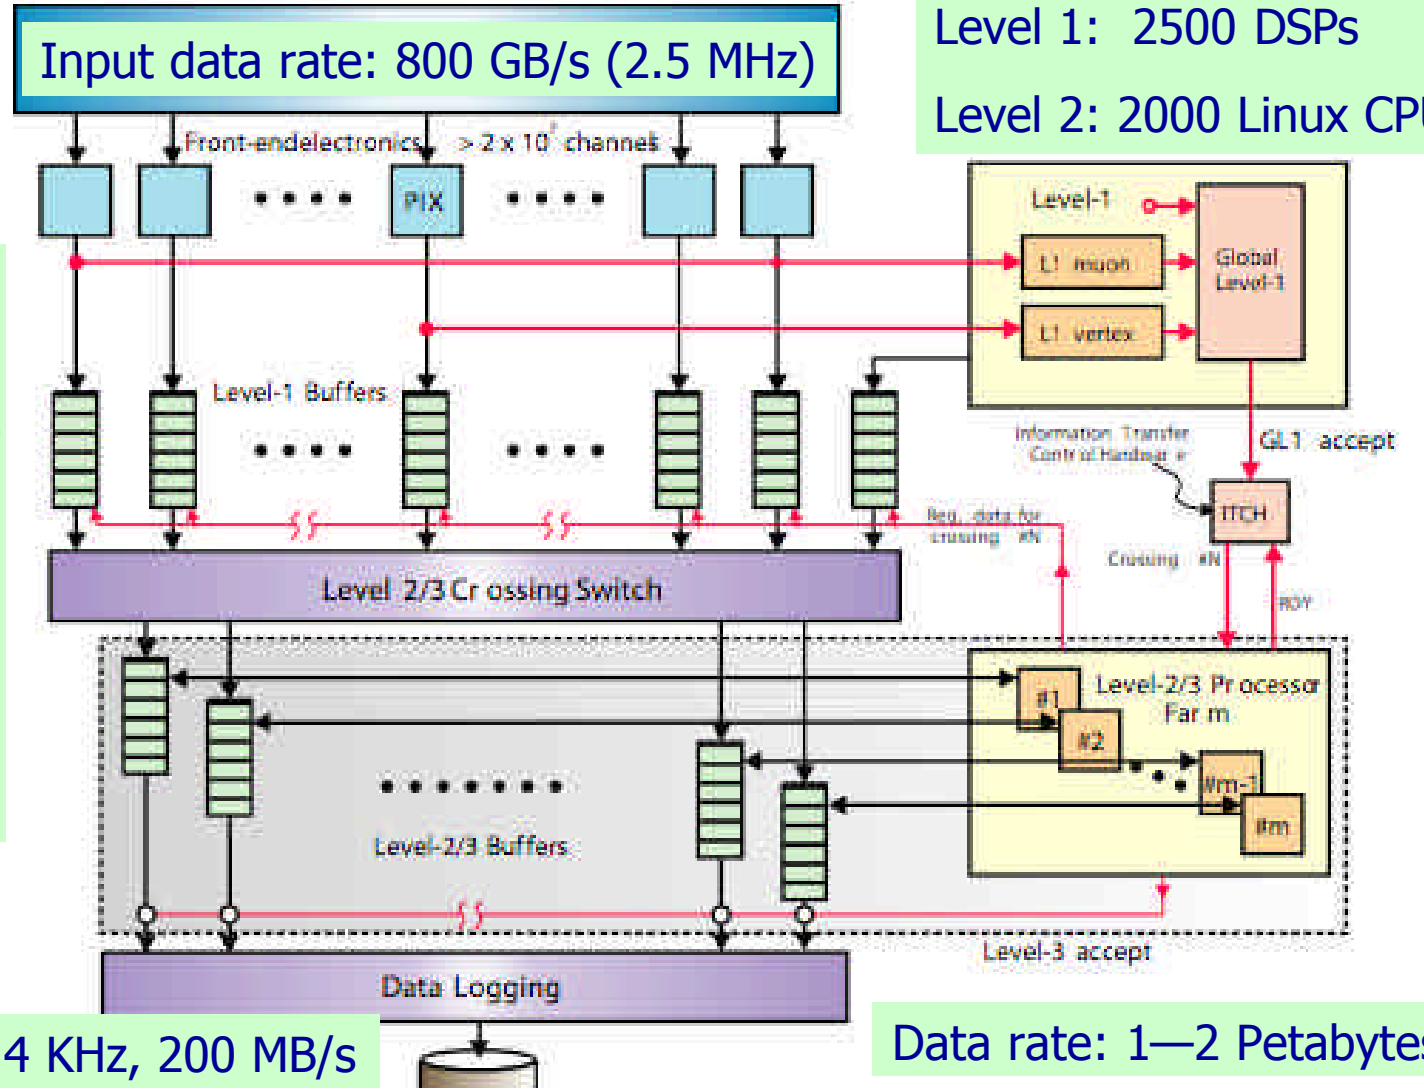

Level 1: 2500 DSPs

Level 2: 2000 Linux CPUs

Pipelined w/ 1 TB buffer, no fixed latency

Output: 4 KHz, 200 MB/s

Data rate: 1—2 Petabytes/yr

The Problem

- The BTeV trigger has very large number of detector electronics and computing resources
  - ~2500 embedded processors for level 1
  - ~2000 PCs for level 2/3
  - ~25,000,000 detector channels
  - Millions of lines of code
- Real-time operation w/ no fixed time latency, averaging:
  - 300us for level-1 decision
  - 13ms for level-2 for decision
  - 130ms level-3 decision
- Failures happen a few times a week for commodity parts
- Software reliability depends on
  - Detector-machine performance
  - Program test procedures, implementation, and design quality
  - Behavior of the electronics (front-end and within trigger)

Fault Adaptation in BTeV

- Implement a large, aggressive trigger, that
  - Applies computation to every interaction
  - Has high sustained computational performance
  - Maintains functional integrity for long periods of time
  - Is highly available
  - Is dynamically reconfigureable, maintainable, and evolvable
- Create fault handling infrastructure capable of
  - Accurately identifying problems (where, what, and why)
  - Compensating for problems (shift the load, changing thresholds)
  - Automated recovery procedures (restart / reconfiguration)
  - Accurate accounting
  - Being extended (capturing new detection/recovery procedures)
  - Policy driven monitoring and control
- Simplify operations

Fault Adaptive Real Time Systems R&D

- These problems are the subject of significant activity in Computer Science and Engineering, but this activity deals with smaller systems and portions of the full problem. The size and scale of the BTeV application is unique and very interesting, and allows investigators to extend and integrate their ideas to a very large, fully functional system.
- A match made in heaven! Both sides have what the other wants.
- Collaboration: "BTeV RTES". Funded by $5M NSF ITR.



Illinois **Pittsburgh** **Syracuse** **Vanderbilt** **Fermilab** **NSF**

How are we attacking the problem?

- **Modeling and Evaluation Framework**: Vanderbilt (with input from Syracuse and Pittsburgh) in the partitioning, load balancing and task allocation parts
- **Runtime Fault Tolerant System**: Illinois, Pittsburgh, and Syracuse, combining VLAs and ARMORs to create the system hierarchy
- **Interface to BTeV, Run Control & Monitoring**: Fermilab
- **Trigger Algorithms, Physics Apps, Input on Operating Conditions**: BTeV Physicists

Hierarchical Approach

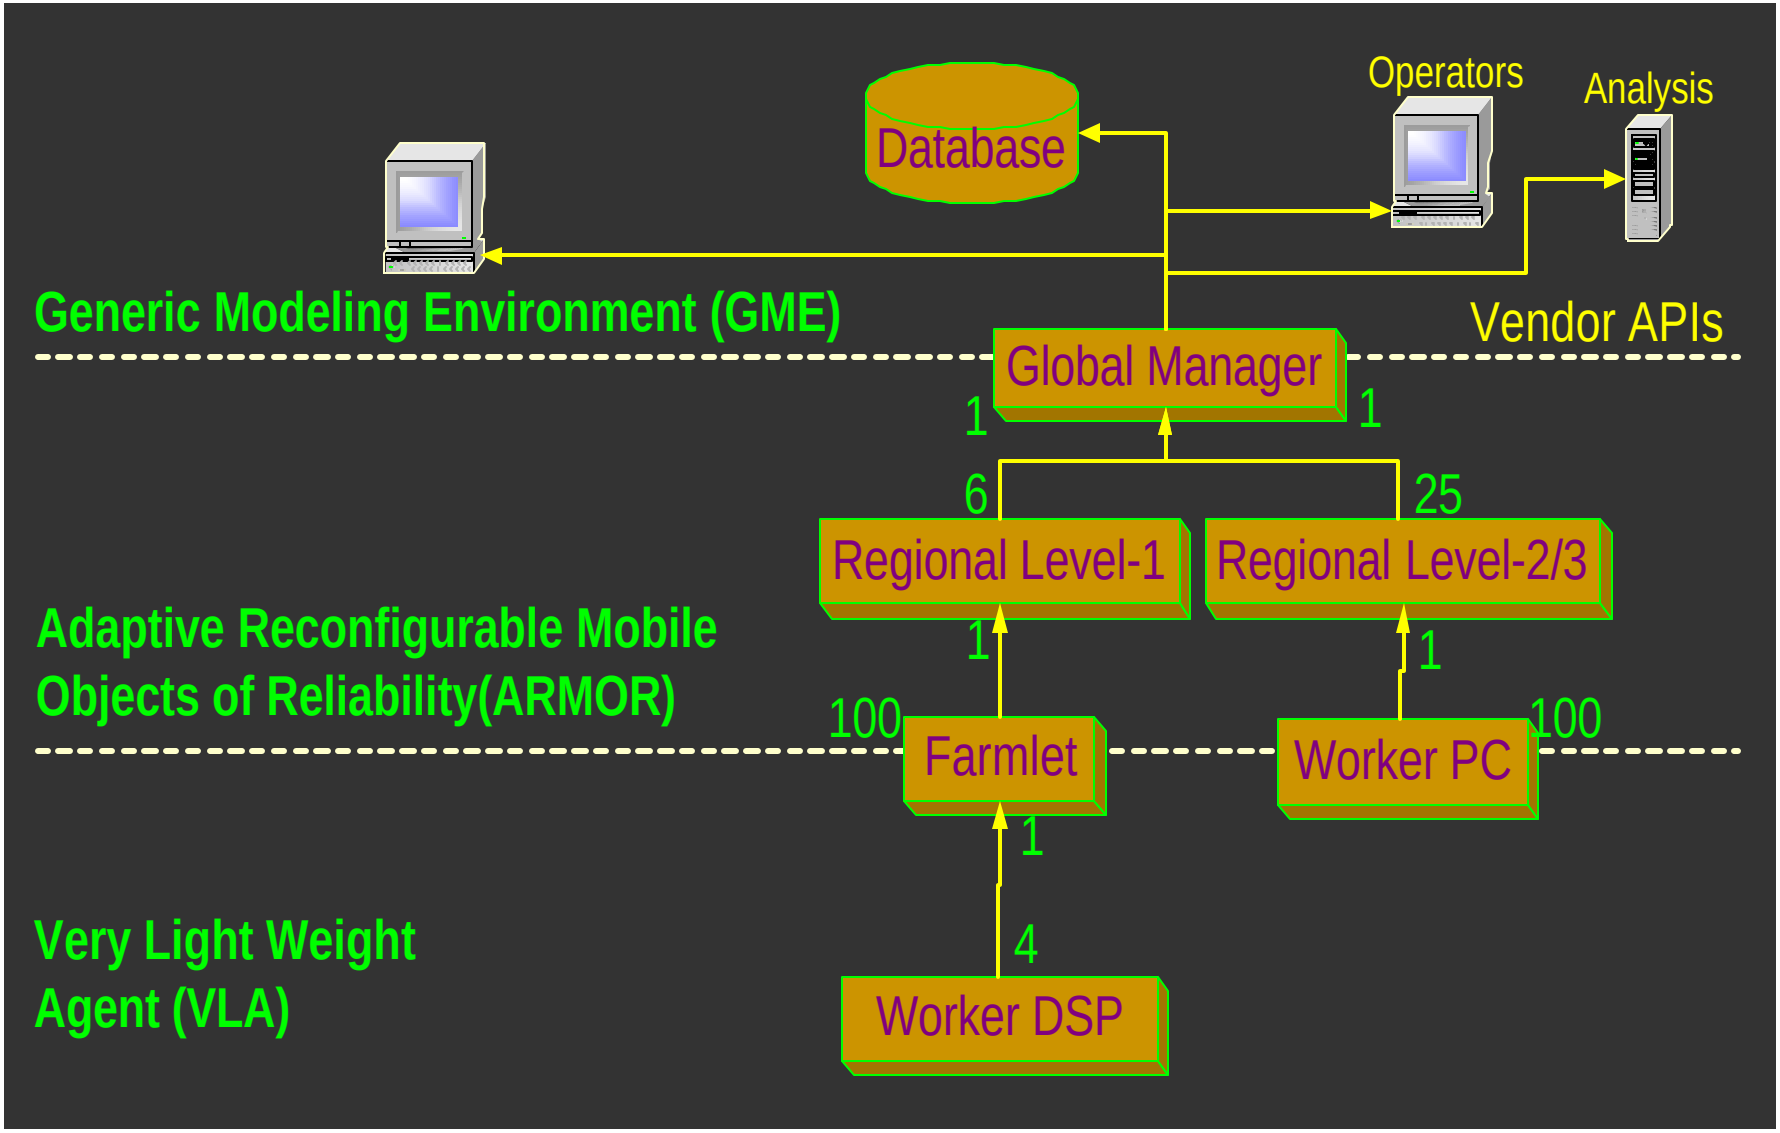

Very Lightweight Agents (VLA)

- Message scheduling & priority assignments
- Fast, simple reactive decisions
- Reads, summarizes, & reports sensors data
- Are "*pluggable*" components
- Lives alongside application
- Some predictive capabilities

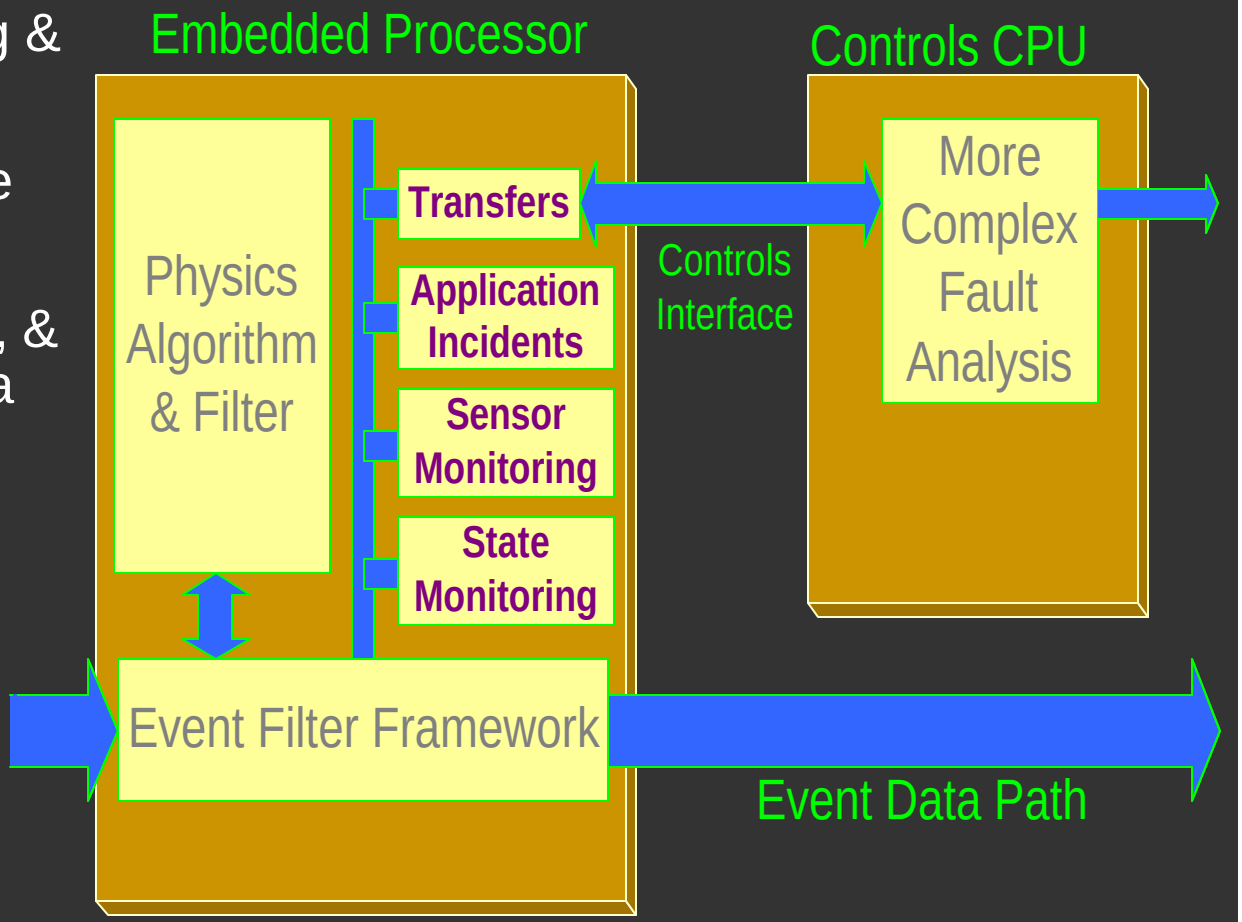

ARMORs

- Are multithreaded processes composed of replaceable building blocks called Elements
- Provide error detection and recovery services to the trigger and other applications
  - Restarts, reconfiguration
  - Removal from service
- A Hierarchy of ARMOR processes form a reconfigurable runtime environment:
  - System management, error detection, and error recovery services are distributed across ARMOR processes
  - ARMOR runtime environment can handle self failure
- ARMOR support for the application:
  - Completely transparent and external support
  - Instrumentation with ARMOR API

Why is all of this interesting?

- It is an integrated approach – from hardware to physics algorithms
  - Standardization of resource monitoring, management, error reporting, and integration of recovery procedures can make operating the system more efficient and make it possible to comprehend and extend.
- There are real-time constraints
  - Scheduling and deadlines
  - Numerous detection and recovery actions
- The product of this research will
  - Automatically handle simple problems that occur frequently
  - Be as smart as the detection/recovery modules plugged into it
- The product can lead to better or increased
  - Trigger uptime by compensating for problems or predicting them instead of pausing or stopping a run
  - Resource utilization - the trigger will use resources that it needs
  - Understanding of the operating characteristics of the software
  - Ability to debug and diagnose difficult problems

Final Thoughts

- This is a highly scalable approach
  - Actions taken as close to the problem as reasonably possible
  - New detection/action elements can be dynamically added to the system
- New and valuable experiences and software are available for use in the BTeV trigger
- Research and development meant to be widely applicable
- This project is a collaboration of physicists and computer scientists, perhaps a model of what it takes to make progress in advanced computing and large scale systems.

Summary

Issues for Science and Engineering Researchers in the Digital Age

What Will This New World Look Like?

- Development is being driven by applications
- Cross-disciplinary teams and efforts are forging solutions
- World will be built from the ground-up… by people on the front lines
- The best researchers will find ways to mold and adapt the developing cyberinfrastructure to break new ground in addressing the most important questions in their fields.

Backup Slides

Fueling Discovery

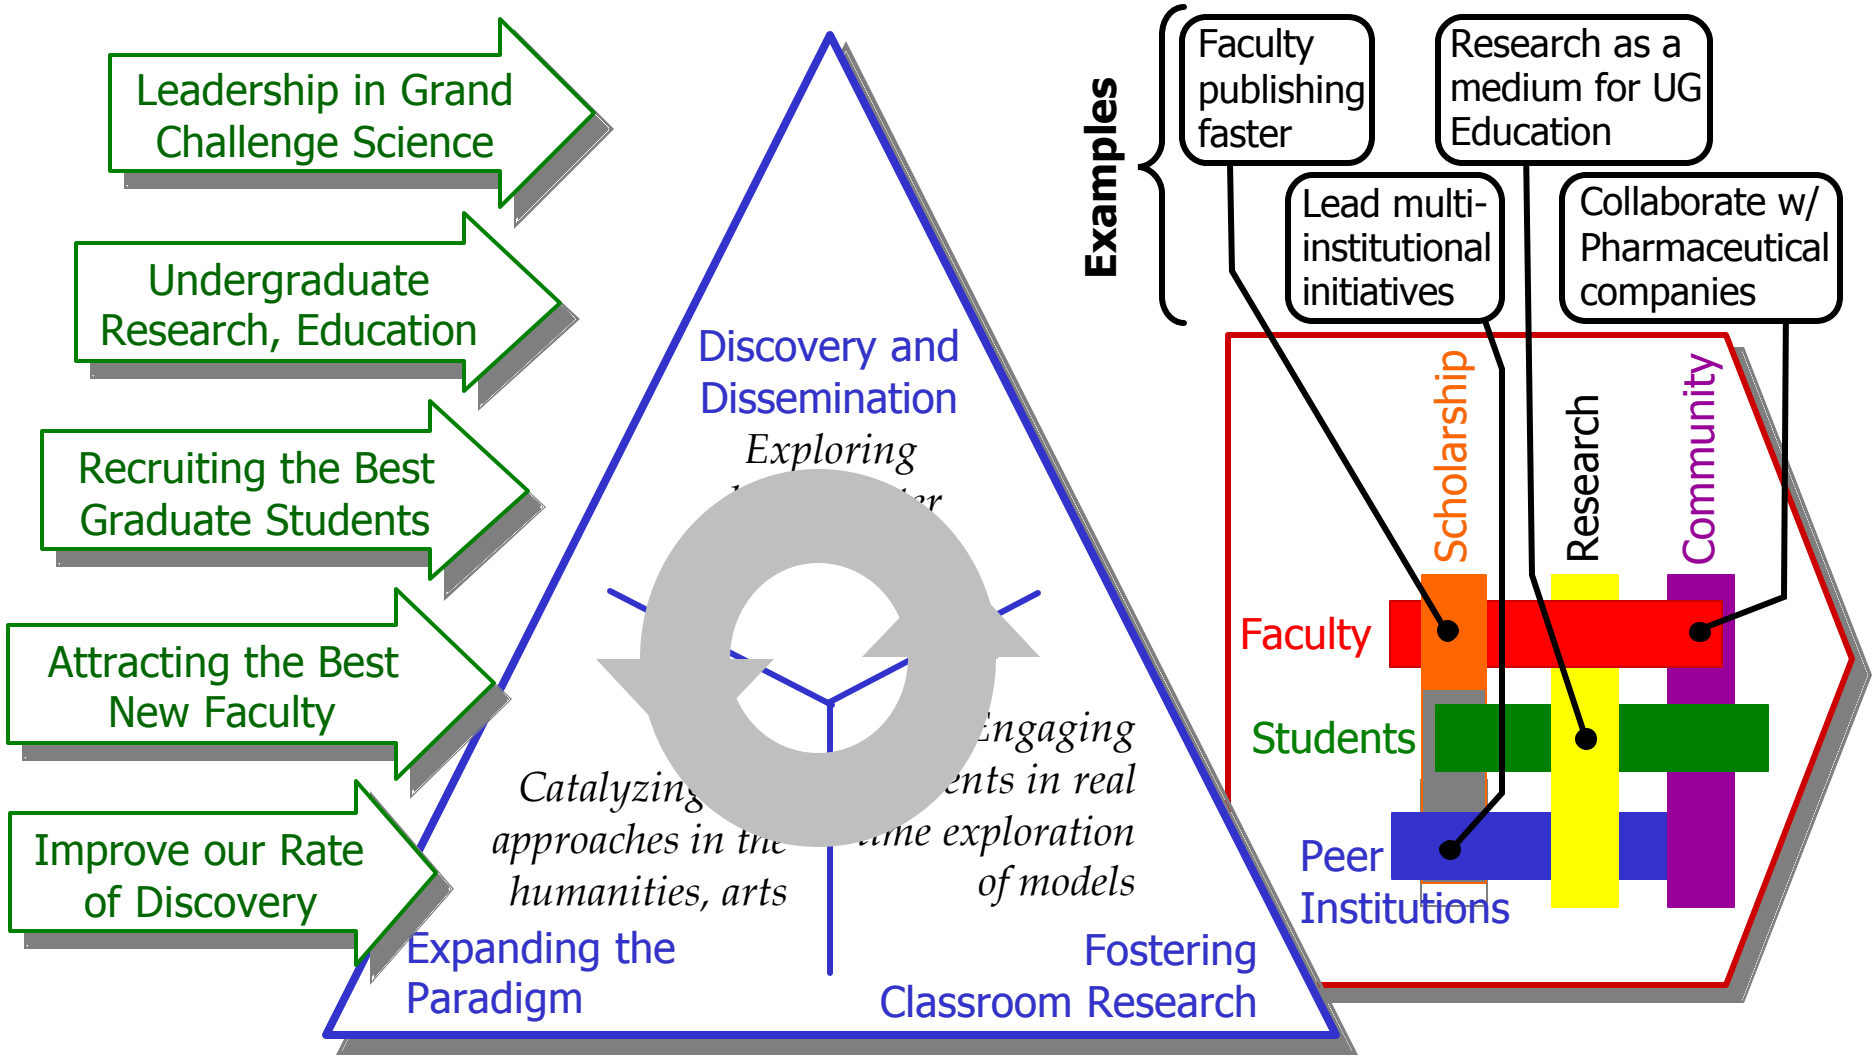

Motivations.....Center Activities.....Results

Storage

- High end storage with lots of redundancy
  - EMC CX600
- Commodity Storage
  - EonStor
- Near-line tape storage
  - Quantum P7000, PX720